



(19)

Europäisches Patentamt
European Patent Office
Office européen des brevets



(11)

EP 0 952 753 A2

(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:
27.10.1999 Bulletin 1999/43

(51) Int. Cl.⁶: H04Q 11/04, H04L 12/56

(21) Application number: 99106051.8

(22) Date of filing: 25.03.1999

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE
Designated Extension States:
AL LT LV MK RO SI

(30) Priority: 30.03.1998 CA 2233491
11.12.1998 US 208984

(71) Applicant:
Newbridge Networks Corporation
Kanata, Ontario K2K 2E6 (CA)

(72) Inventor: Dubuc, Ken
Ottawa, Ontario K2B 8K3 (CA)

(74) Representative:
Prato, Roberto et al
STUDIO TORTA S.r.l.,
Via Viotti 9
10121 Torino (IT)

(54) ATM adaption layer traffic scheduling for variable bit rate connections

(57) The ATM apparatus for transmitting messages associated with a plurality of variable bit rate connections comprises ATM layer transmission equipment for receiving multiple messages, segmenting each received message into a plurality of smaller data units, such AAL SAR PDUs, and multiplexing such data units into a single stream for transport over a physical interface. A scheduler receives messages from each of the variable bit rate connections and transfers the messages to the transmission equipment in an order corresponding to the level of a leaky bucket associated with each connection. The scheduler also dynamically sets a transmission rate, 1/T, for each connection at the time its message is transferred to the transmission equipment, where

$$T \leftarrow \max\left(T_s + \frac{X - \tau_s}{N - 1}, T_p, T_L\right), \quad \text{if } N > 1$$

connections and minimize interleaving of data units from the various connections in order to reduce the average latency (per unit of message) in transmitting the messages across a network.




Fig. 1

and

$$T \leftarrow \max(T_p, T_L), \quad \text{if } N = 1$$

where T_s is a period corresponding to a constant sustained transmission rate, T_p is a period corresponding to a peak transmission rate, τ_s is a burst tolerance, N is the number of data units in the message, X is a fill level of the leaky bucket associated with the given connection, and T_L corresponds to all unused bandwidth allocated to a service class to which the given connection belongs. These actions increase the burstiness of the




Fig. 5

Description**Field of Invention**

- 5 [0001] The invention relates generally to the art of digital communications and more specifically to a system for minimizing the average latency in transporting messages, such as packets or frames, which are segmented into a plurality of smaller cells for transport across a network.

Background of Invention

- 10 [0002] Asynchronous transfer mode (hereinafter "ATM") service inter-networking protocols enable data or messages formatted according to a non-ATM data communication protocol to be transported across an ATM network. For example, the Frame Relay Forum FRF.5 protocol specifies how a relatively large, variable length, frame relay packet should be segmented into a plurality of ATM-like, fixed-size, cells for transport across an ATM network. Such protocols necessarily define how the ATM Adaption Layer (AAL) should be provisioned since this layer of the ATM/B-ISDN protocol stack, as defined by ITU Recommendation I.321 and shown in Fig. 1, is responsible for adapting the services provided by the ATM Layer, which provides basic ATM cell transport functions, to higher layers, e.g. frame relay bearer service.
- 15 [0003] Fig. 2 illustrates a generic version of the AAL in greater detail. As shown in Fig. 2, some versions of the AAL, such as AAL3/4 and AAL5, include a convergence sublayer (CS) and a segmentation and reassembly sublayer (SAR).
- 20 [0004] The CS, which sits directly above the SAR and below the AAL Service Access Point (SAP), aligns the SDUs and adds overhead information. The CS may also provide service specific signalling or data link functions.
- [0004] The SAR, when operating in a message mode, segments a single AAL SAR Service Data Unit (hereinafter "SDU"), such as a variable length frame packet, into a plurality of AAL SAR Protocol Data Units (hereinafter "PDU"), each of which essentially forms the payload of an ATM cell transmitted across an ATM network. Conversely, at the destination, the destination SAR requires that all of the PDUs composing an SDU be passed from the ATM Layer to the destination SAR before it can reassemble the SDU and, ignoring the role of the convergence sublayer, indicate reception of the SDU to the higher layer using the AAL. Thus, the latency in transmitting an SDU from a first point to a second point in a network can be defined as the time from which the transmission of the SDU is first requested until the time the last PDU arrives at the destination SAR and the SDU is reassembled. In other words, latency can be defined as the time required to transmit the SDU from an originating AAL SAP to a destination AAL SAP. This latency is entirely characterized by the amount of time required to propagate the last PDU of an SDU across the ATM network - the time required to propagate any other PDU before the last PDU of an SDU is of no consequence at the destination AAL SAP.
- 25 [0005] Latency manifests itself as sluggishness or slow response time in interactive-type communications. For example, if one were sending joystick instructions across a network during the course of an interactive game played there over, a long latency would, in the absence of other aggravating factors, result in a noticeable time period between the physical movement of the joystick and the corresponding computer action. Accordingly, it is desirable to minimize latency for interactive telecommunications applications.
- 30 [0006] Latency is affected by the service discipline used to schedule or multiplex PDUs corresponding to SDUs from a plurality of virtual connections (VCs) into a single cell stream for transmission across the Physical Layer (PHY) of the ATM network. Fig. 3 shows how an ATM Layer 11 provides a SAP 10 to each of several VCs, each of which has its own AAL 12 (i.e., the AAL is invoked in parallel instances). The ATM Layer 11, in turn, uses a single SAP 14 into a PHY 16. One role of the ATM Layer 11 is to accept requests of PDUs 17 from each SAP 10 and to multiplex these PDUs into a single cell stream 18 such that the timing of the transmission of each of the PDUs conforms to predetermined traffic parameters assigned to its respective VC.
- 35 [0007] Fig. 3 illustrates a condition where each VC generates a burst 20 of several ATM PDU requests at the ATM SAP 10, wherein each such burst corresponds to a single SDU 22, such that there is an overlap in the transmission periods of the SAR SDUs from the ATM Layer 11 to the PHY 16. The ATM PDUs 17 received from the ATM SAP 10 must therefore be queued, and then the ATM PDUs from each of the different ATM SAPs 10 must be multiplexed in some order onto the single stream 18 of ATM PDUs passed to the PHY SAP 14. Given this set of PDUs which have been requested over several ATM Layer SAPs 10, and subject to the constraints of satisfying the traffic parameters of each VC, it is often desired to minimize the average amount of latency experienced per unit of SDU data (i.e., per PDU) for various types of ATM service categories.
- 40 [0008] As shown in Fig. 3, a typical ATM Layer implementation might use round-robin ordering in sending the PDUs 17 to the PHY SAP 14 from each ATM SAP 10. This would result in each corresponding SDU 22 using an equal fraction of the PHY bandwidth while the PDUs for each SDU are being transmitted. This is shown, for instance, in the bandwidth occupancy chart of Fig. 4A for the situation where two VCs each request a burst of the same number of PDUs at about the same time, wherein each VC has a PCR equal to 100% of the available bandwidth. (A "bandwidth occupancy chart" is a chart with time on the horizontal axis, and bandwidth on the vertical axis. Each SDU sent on an ATM virtual con-

nection is shown as a shaded region on such a chart. The net height of the region at a particular time shows the amount of bandwidth occupied by the transmission of the SDU at that time; the leftmost and rightmost extent of the region gives the time at which the first and last PDUs for the SDU are transmitted, respectively; and the total area of the region gives the size of the SDU. Unshaded regions in these charts represent the proportion of unused PHY bandwidth, for which the ATM Layer will be sending idle cells.) This ordering is not optimal with respect to the average amount of latency experienced per unit of SDU data.

Summary of Invention

- 10 [0009] Broadly speaking, the invention seeks to minimize or reduce the average per unit latency in transporting messages which are decomposed into a plurality of smaller data units for transport across a network.
- [0010] One aspect of the invention relates to an apparatus for transmitting messages associated with a plurality of variable bit rate connections, each of which is associated with a traffic contract which defines compliance thereto as conformance to a leaky bucket algorithm. The apparatus comprises transmission equipment for receiving multiple messages, segmenting each received message into one or more data units, and multiplexing such data units from various connections into a single stream for transport over a physical interface at an output transmission rate. A bandwidth provisioning means associated with the transmission equipment dynamically provides a portion of the output transmission rate to any connection. A scheduler schedules the transfer of messages to the single stream of the transmission equipment and directs the bandwidth provisioning means to allocate a portion of the output transmission rate to each connection at the time its message is transferred to the transmission equipment. The portion of the output transmission rate for a given connection is substantially equal to $1/T$, T being computed as

$$T \leftarrow \max\left(T_S + \frac{X \cdot \tau_S}{N - 1}, T_p, T_L\right), \quad \text{if } N > 1$$

25

and

$$T \leftarrow \max(T_p, T_L), \quad \text{if } N = 1$$

30

where T_S is a period corresponding to a constant sustained transmission rate, T_p is a period corresponding to a peak transmission rate, τ_S is a burst tolerance, N is the number of data units in the message, X is a fill level of the leaky bucket associated with the given connection, and T_L corresponds to all unused bandwidth allocated to a service class to which the given connection belongs.

- 35 [0011] Another aspect of the invention relates to an apparatus for transmitting messages associated with a plurality of variable bit rate connections, each of which is associated with a traffic contract which defines compliance thereto as conformance to a leaky bucket algorithm. The apparatus comprises transmission equipment for receiving multiple messages, segmenting each received message into one or more data units, and multiplexing such data units from various connections into a single stream for transport over a physical interface at an output transmission rate. Bandwidth provisioning means associated with the transmission equipment dynamically provides a portion of the output transmission rate to any connection. A scheduler schedules the transfer of messages to the single stream of the transmission equipment. The scheduler is enabled to transfer the messages according to an order corresponding to the level of a leaky bucket associated with each corresponding connection in order to reduce the interleaving of data units from various connections over the single stream.

- 40 [0012] In the preferred embodiment, the transmission equipment comprises a segmentation means and a multiplexing means for respectively segmenting each received message into data units and multiplexing the data units into the aforesaid single stream. The scheduler preferably transfers messages to the single stream of the transmission equipment by receiving the messages from higher layer networking services and transferring the messages to the segmentation means which substantially immediately transfers the messages to the multiplexing means. Alternatively, the scheduler may transfer messages to the single stream of the transmission equipment by receiving the data units from the segmentation means and transferring groups of data units to the multiplexing means, wherein each such group corresponds to a message.

- 45 [0013] In the preferred embodiment, the transmission equipment includes an indication means for informing the scheduler when a message has been fully transmitted by the transmission equipment. When the indication means is actuated, the scheduler is enabled to set a new transmission rate for a given connection during the time a message associated therewith is in the process of being transmitted by the transmission equipment. The new transmission rate is substantially equal to $1/T$, T being computed as

$$T \leftarrow \max\left(T_S + \frac{X_{est} - \tau_S}{R_{est} - 1}, T_P, T_L\right), \quad \text{if } R_{est} > 1,$$

5 where X_{est} and R_{est} respectively represent an expected bucket fill level and an estimated number of data units remaining to be transmitted in the message for the given connection.

[0014] The most preferred embodiment relates to an ATM communications device wherein the messages are ATM adaption layer (AAL) service data units (SDUs) associated with a plurality of virtual circuit connections, and the data units are AAL segmentation and reassembly protocol data units (SAR PDUs). The transmission equipment preferably 10 comprises segmentation means for segmenting each AAL SDU into one or more SAR PDUs, ATM layer means for multiplexing the SAR PDUs of various connections into a single stream having an output transmission rate, and physical transmission means, connected to the ATM layer means, for transporting the single stream of SAR PDUs over a physical interface.

[0015] Another aspect of the invention relates to a method for transmitting multiple messages associated with multiple 15 variable bit rate connections over a network, wherein each such connection is associated with a traffic contract which defines compliance thereto as conformance with a leaky bucket algorithm. The method comprises the steps of:

- (a) segmenting each message of each connection into one or more transportable data units;
- 20 (b) multiplexing the data units from various connections into a single stream for transport over a transmission media to a destination;
- (c) scheduling the transfer of messages from the multiple connections to the multiplexing step; and
- 25 (d) allocating a portion of the transmission media bandwidth to each connection substantially when its message is transferred to the multiplexing step, the portion of the transmission media bandwidth for a given connection being substantially equal to $1/T$, T being computed as

30

$$T \leftarrow \max\left(T_S + \frac{X \cdot \tau_S}{N - 1}, T_P, T_L\right), \quad \text{if } N > 1$$

35 and

$$T \leftarrow \max(T_P, T_L), \quad \text{if } N = 1$$

where T_S is a period corresponding to a constant sustained transmission rate, T_P is a period corresponding to a peak 40 transmission rate, τ_S is a burst tolerance, N is the number of data units in the message, X is a fill level of the leaky bucket associated with the given connection, and T_L corresponds to all unused bandwidth allocated to a service class to which the given connection belongs.

[0016] Another aspect of the invention relates to a method for transmitting multiple messages associated with multiple 45 variable bit rate connections over a network, wherein each such connection is associated with a traffic contract which defines compliance thereto as conformance with a leaky bucket algorithm. The method comprises the steps of:

- (a) segmenting each message of each connection into transportable data units;
- 50 (b) multiplexing the data units from various connections into a single stream for transport over a transmission media to a destination; and
- (c) scheduling the transfer of messages from the multiple connections to the multiplexing step according to an order corresponding to the level of a leaky bucket associated with each corresponding connection in order to reduce the interleaving of data units from various connections over the single stream.

55

Brief Description of Drawings

[0017] For the purposes of description, but not of limitation, the foregoing and other aspects of the invention are

explained in greater detail with reference to the accompanying drawings, wherein:

- Fig. 1 is an illustration of an ATM protocol stack;
- Fig. 2 is a protocol diagram illustrating the sub-layers of an ATM Adaption Layer (AAL) shown in Fig. 1;
- Fig. 3 is a data flow diagram illustrating the flow of PDUs through an ATM Layer shown in Fig. 1 using a round-robin scheduling discipline for servicing parallel instances or invocations of the AAL;
- Fig. 4A is a bandwidth occupancy chart illustrating the average latency in the transmission of SDUs from two virtual circuits using the round-robin scheduling discipline depicted in Fig. 3;
- Fig. 4B is a bandwidth occupancy chart exemplifying the average latency in the transmission of SDUs from two virtual circuits using apparatus according to one aspect of the invention;
- Fig. 5 is a protocol diagram illustrating the structure of the AAL in accordance with a first preferred embodiment of the invention which includes an SDU scheduling sublayer;
- Figs. 6A and 6B are flowcharts illustrating a first preferred method, carried out by the scheduling sublayer, for servicing SDU data requests;
- Fig. 7 is a diagram illustrating various queues employed by the scheduling sublayer according to the first preferred embodiment;
- Fig. 8 is a table exemplifying average latency in transmitting SDUs using the first preferred method and using a prior art round-robin multiplexing scheme;
- Fig. 9 is a bandwidth occupancy chart corresponding to the example in Fig. 8 where SDUs are transmitted using the first preferred method;
- Fig. 10 is a bandwidth occupancy chart corresponding to the example in Fig. 8 where SDUs are transmitted using the prior art round-robin multiplexing scheme;
- Fig. 11 is a flow chart of a leaky bucket algorithm;
- Fig. 12 is a bandwidth occupancy chart exemplifying a transmission pattern for non-real-time VBR VCs in accordance with the first preferred method;
- Fig. 13 is a bandwidth occupancy chart exemplifying a transmission pattern for the same non-real-time VBR VCs shown in Fig. 12 but in accordance with a second preferred method for scheduling SDUs;
- Fig. 14A and 14B are flowcharts illustrating the second preferred method, carried out by the scheduling sublayer, for servicing SDU data requests;
- Fig. 15 is a bandwidth occupancy chart illustrating a range of possible transmission times for PDUs of an SDU associated with a non-real-time VBR VC;
- Fig. 16 is a bandwidth occupancy chart exemplifying a transmission pattern for non-real-time VBR VCs in accordance with the second preferred method; and
- Fig. 17 is a bandwidth occupancy chart exemplifying a transmission pattern for the same non-real-time VBR VCs shown Fig. 16 but in accordance with the first preferred method.

Detailed Description of Preferred Embodiments

[0018] The detailed description is divided into two parts. The first part discloses a preferred embodiment for minimizing the average latency per unit of AAL SDU data in relation to ATM service classes such as the available bit rate (ABR) and unspecified bit rate (UBR) service categories which are characterized by a peak cell rate (PCR) traffic management parameter. The second part discloses a preferred embodiment for minimizing the average latency per unit of AAL SDU data in relation to ATM service classes such as non-real-time VBR which are characterized by peak cell rate (PCR), sustained cell rate (SCR), and maximum burst size (MBS) traffic management parameters. As will become apparent, the second embodiment builds upon the first embodiment.

Service Classes Characterized by PCR

[0019] Fig. 5 illustrates the structure of an AAL in accordance with the preferred embodiment. In addition to the known SAR layer 12 and CS layer 13 (which is not explicitly shown in Fig. 3), the preferred AAL includes a scheduling means, such as scheduling sublayer 24 (alternatively referred to as "SS"), located above the CS/SAR layers for scheduling or ordering the transfer of SDUs to the CS/SAR. (In alternative embodiments, scheduling sublayer 24 may be placed above the SAR but below the CS.)

[0020] The scheduling sublayer 24 provides a plurality of SS SAPs 26, one for each VC associated with PHY SAP 14. The scheduling sublayer 24 accepts AAL SDUs 22 in data requests from higher layers of the ATM protocol stack (hereinafter "AAL user"). In the preferred embodiment, the AAL user uses the same data request to the scheduling sublayer at SS SAPs 26 as the AAL user would in the prior art for direct access to the AAL SAR SAP 15. The scheduling sublayer 24, in turn, preferably communicates with the SAR 12 using prior art AAL SAR SAP 15. As will be described in greater

detail below, the scheduling sublayer 24 affects only the egress data path (from AAL user to PHY) of SDUs 22, but not the ingress data path (from PHY to AAL user).

[0021] The scheduling sublayer 24 also communicates directly with the ATM Layer 11 through a cell rate specification signal 28 and an SDU transmission signal 30. The cell rate signal 28 enables the scheduling sublayer 24 to specify a constant cell rate or period for any given VC to conventional ATM Layer equipment 11' which, as known *per se*, provides this capability. The SDU transmission signal 30 enables the ATM Layer equipment 11' to inform (such as by way of a software interrupt) the scheduling sublayer 24 when a last PDU of any given SDU 22 and VC has been transmitted. This capability is also provided by conventional ATM Layer equipment 11', as known in the art *per se*, based upon examination of certain control bits in the PDU or when the equipment no longer has any cells to send for a particular connection.

[0022] The scheduling sublayer 24 is provided with data regarding the traffic parameters, such as the PCR, and current transmit state of each VC; information which is typically known by the ATM Layer but not the CS/SAR. The scheduling sublayer 24 is also provided with data regarding SDU boundaries; information which is typically known by the CS/SAR but not the ATM Layer. This data is collectively used by the scheduling sublayer 24 to control the time at which SDUs 22 are passed to the CS/SAR, and to specify a constant cell rate for a given VC at the ATM Layer 11.

[0023] The scheduling sublayer 24 attempts to minimize the effects of the ATM Layer 11 which interleaves PDUs associated with various VCs in round-robin fashion for transport through PHY SAP 14. This is accomplished by transferring the AAL SDUs of VCs to the CS/SAR layers on a substantially relative first-come, first-served basis (of VCs making data requests at SS SAPs 26) subject to the constraints that: (a) the PHY bandwidth consumed by the VCs are maximized within the limits of the bandwidth preallocated to their respective service classes; (b) the traffic parameters for each VC are satisfied; and (c) the SDUs, and the PDUs thereof, of any given VC be transmitted across PHY SAP 14 in sequential order so as to prevent misinsertion at the destination. This approach attempts to minimize the average transmission latency encountered by the AAL SDUs 22.

[0024] The potential results in accordance with this approach are simplistically illustrated in Fig. 4B for the situation, corresponding to Fig. 4A, wherein two VCs each request the transmission of one SDU having the same number of PDUs at about the same time, with each VC having a PCR equal to 100% of the available bandwidth. As shown in Fig. 4B, each VC is given a maximal amount of bandwidth, i.e., up to its PCR, but the average latency is reduced compared to Fig. 4A.

[0025] Figs. 6A and 6B illustrate in flow chart form the preferred scheduling method carried out by the scheduling sublayer 24. Fig. 7 illustrates various queues employed by the scheduling sublayer 24. A TX list or queue 34 stores a list of VCs which are currently being transmitted by the CS/SAR and ATM Layer; that is, those VCs which have had an SDU (i.e., all of the PDUs thereof) passed to SAR 12 by the scheduling sublayer 24. Queue 34 also stores the current cell rate for each VC listed therein. A TX-pending queue 36 stores a list of all VCs which have made data requests to the scheduling sublayer 24, but the data requests and SDUs thereof have not yet been passed by the scheduling sublayer to AAL SAR SAP 15. In addition, for each VC listed in queue 36, there exists one VC buffer or queue 38A, 38B, 38C, 38D, ... or 38N for storing unprocessed SDUs of the corresponding VC.

[0026] There are two triggering events in the preferred method. Event 40 (Fig. 6A) is the occurrence of a data request from the AAL user at any SS SAP 26. Event 60 (Fig. 6B) is the communication of signal 30 by ATM Layer equipment 11' to the scheduling sublayer 24 indicating that the equipment has finished transmitting an SDU (i.e., the last PDU thereof has been transmitted) for a given VC.

[0027] Referring to Figs. 5, 6A and 7, when the scheduling sublayer 24 receives at any SS SAP 26 a data request and the SDU thereof for a given VC, VC(i), the scheduling sublayer 24 checks at step 42 whether or not VC(i) is currently being serviced by it, i.e., if VC(i) is in the TX queue 34 or TX-pending queue 36. If so, at step 46, the scheduling sublayer 24 stores the SDU in a corresponding VC_i queue 38. These steps ensure that the AAL SDUs of a given VC are transferred to the CS/SAR in sequential order.

[0028] If the condition at step 42 is not satisfied, then a check is made at step 44 whether or not there is any unused PHY bandwidth for the service class that VC(i) belongs to. If PHY bandwidth is available, then at steps 48, 50, 52 and 54 the scheduling sublayer 24:

(a) instructs the ATM Layer to set the current cell rate for VC(i) to the lower of: (i) the PCR of VC(i), and (ii) the remaining unused PHY bandwidth for the service class that VC(i) belongs to;

(b) decreases the available PHY bandwidth for the service class that VC(i) belongs to by its current cell rate;

(c) passes the SDU to the CS/SAR which, practically speaking, very quickly completes the segmentation of the SDU into PDUs and passes these to the ATM Layer for transmission; and

(d) adds VC(i) to the tail end of the TX queue 34.

[0029] If, however, no PHY bandwidth is available at step 44, at step 56 VC(i) is added to the tail end of the TX-pending queue 36 and the SDU of VC(i) is added to the tail end of the corresponding VC_i queue 38.

[0030] Referring to Figs. 5, 6B and 7, when the scheduling sublayer 24 receives an indication at event 60 that the ATM Layer 11 has finished sending all PDUs corresponding to an SDU of a given VC, VC(k), through the PHY SAP 14, the SS 24 proceeds at step 62 to remove VC(k) from the TX queue 34 and update the PHY bandwidth availability for the service class that VC(k) belongs to. At step 64, a check is made whether or not VC(k) has any AAL SDUs pending transmission, that is, if any SDUs are present in VC_k queue 38. If so, at step 66 VC(k) is placed on the TX -pending queue 36.

[0031] Steps 68, 72 and 74 set up a loop for scanning the TX queue 34, starting from its head. Step 70 determines whether the current cell rate of any VC, VC(x), in queue 34 can be increased. If so, at step 76 the scheduling sublayer 24 signals ATM Layer 11 to increase the current cell rate for VC(x) to the lower of: (i) the PCR of VC(x), and (ii) the remaining unused PHY bandwidth for the ATM service class that VC(x) belongs to. The available PHY bandwidth for that service class is also decreased accordingly. The TX queue 34 is scanned until the end of the queue or until the total PHY bandwidth available at PHY SAP 14 is exhausted.

[0032] Steps 78, 82 and 84 set up a loop for sequentially scanning the TX-pending queue 36, starting from its head. If at step 80 PHY bandwidth remains available in the service class of a given VC, VC(x), at the head of queue 36, then at step 86 the scheduling sublayer 24 removes VC(x) from TX-pending queue 36. The scheduling sublayer 24 also removes the AAL SDU at the head of the corresponding VC_x queue 38, and processes the removed SDU as if a data request for VC(x) was received at the SS SAP 26, as per steps 48 through 58 discussed above. The TX-pending queue 36 is scanned until the end of the queue or until the total PHY bandwidth available at PHY SAP 14 is exhausted.

[0033] Referring to Figs. 8, 9 and 10, an example of how the scheduling sublayer 24 provides a lower average latency in transmitting SDUs than the prior art system shown in Fig. 3. Fig. 8 is a table which specifies the VC identifier related to each data request, the time at which the request is made, the number of PDUs in the SDU, and the PCR for the VC. For simplicity, the VCs are assumed to belong to the same service class and the PCR is given as a percentage of the PHY bandwidth. The time has also been unitized so that one time unit is the length of time required for the PHY to send a single cell.

[0034] Fig. 9 is a bandwidth occupancy chart illustrating the latency of each SDU that would likely result from the preferred scheduling method. This latency is the difference between the time a data request for the SDU was first made at SS SAP 26 and the time at which the last PDU for the SDU was sent by the ATM Layer to the PHY. For comparison, Fig. 10 illustrates the latencies which would result under a typical prior art round-robin ATM Layer multiplexing scheme. The average latencies per unit of SDU data, - that is, per PDU - are calculated and tabulated in the Fig. 8 table. The average latencies are calculated by taking the sums of the latency of each SDU multiplied by the number of PDUs in the SDU, and dividing by the total number of PDUs in all SDUs.

[0035] Under the prior art system as shown in Fig. 10, all VCs begin transmission simultaneously, with each VC occupying 25% of the PHY bandwidth. At t = 8, the SDU requested by VC 2 will have completed transmission. Accordingly, the prior art system equally reallocates the PHY bandwidth formerly occupied by VC 2 to VC 1 and VC 4; VC 3 is already transmitting at its PCR and cannot be allocated more bandwidth. The SDUs requested by VC 1 and VC 4 complete transmission at approximately time t = 10.5, (in practice, one at time t = 10 and one at time t = 11). At t = 16, the SDU requested by VC 3 completes transmission. The average latency is computed as $(10.5 * 3 + 8 * 2 + 16 * 4 + 10.5 * 3) / (3 + 2 + 4 + 3) \approx 11.9$.

[0036] Under the preferred method, as shown in Fig. 9, at t = 0, VC 1 is allocated its PCR, and the ATM Layer is instructed to send the PDUs of VC 1 at 50% of the PHY bandwidth. The SDU for VC 1 is passed to the SAR. Similarly, VC 2 is allocated its PCR, and the ATM Layer is instructed to send the PDUs of VC 2 at 33% of the PHY bandwidth, and the SDU for VC 2 is passed to the SAR. VC 3 is allocated the remaining PHY bandwidth, and the SDU for VC 3 is passed to the SAR. VC 1, VC 2 and VC 3 are added to the TX queue 34 of VCs whose SDU transmission is in progress. Finally, since there is no remaining PHY bandwidth for VC 4, the scheduling sublayer adds VC 4 to the TX-pending queue 36 of VCs which have data requests pending. This state is also illustrated in Fig. 7A.

[0037] At t = 6, VC 1 and VC 2 both complete transmission of all PDUs for their SDUs. The scheduling sublayer 24 removes these VCs from the TX queue 34 of VCs whose transmission is in progress, and the PHY is marked as having 83% of its bandwidth previously allocated to those VCs as free. The TX-pending queue 36 of VCs whose transmission is in progress is then scanned, and VC 3 is found to be operating below its PCR, so the ATM Layer is instructed to increase the PHY bandwidth allocated to VC 3 from 17% to 25%. Since there is still free PHY bandwidth, the TX-pending queue 36 of VCs with data requests pending is scanned. VC 4 is found, and the ATM Layer is instructed to allocate 75% of the PHY bandwidth to VC 4 - 75% being both VC 4's PCR and the previously remaining free PHY bandwidth - and the SDU for VC 4 is passed to the SAR. VC 4 is then added to the TX queue 34 of VCs whose transmission is in progress. This state is also illustrated in Fig. 7B.

[0038] At t = 10, VC 4 completes transmission of PDUs for its SDU. VC 4 is therefore removed from the TX queue 34 of VCs whose transmission is in progress. Queue 34 is then scanned. VC 3 is found, but it is already transmitting at its

PCR, so its transmission is not modified. There are no other VCs with data requests pending, so the PHY becomes 75% idle. At $t = 18$, VC 3 completes transmission.

[0039] The average latency under the scheduling sublayer 24 is computed as $(6 * 3 + 6 * 2 + 18 * 4 + 10 * 3) / (3 + 2 + 4 + 3) = 11$.

[0040] The preferred embodiment provides a number of advantages. First, it schedules SDUs to minimize the average amount of latency per unit of AAL SDU data on a given PHY. This results in effectively a lower average delay in data transmission through an ATM network. Second, the preferred method features a low storage and computational load, and can therefore be easily applied in real time systems. Third, the preferred embodiment strives to ensure that all available PHY bandwidth is consumed. Fourth, the preferred embodiment is not dependent on elaborate ATM Layer equipment. The only requirement of the ATM Layer equipment is that it be configurable so that the transmission rate of each connection or VC can be dynamically adjusted.

[0041] However, while the preferred embodiment employs ATM Layer equipment capable of generating a direct indication, e.g. software interrupt, of when an AAL SDU has finished being transmitted through the ATM Layer, alternative embodiments need not provide this capability. Instead, the scheduling sublayer may use a separate timer to provide this indication. This is because the scheduling sublayer 24 is provided with the cell rate of each VC and the number of AAL PDUs to be sent on each connection. Therefore, the scheduling sublayer can compute the approximate time (i.e., ignoring cell delay variation introduced by ATM Layer multiplexing and PHY functions) at which the ATM Layer will finish sending the ATM PDUs to the PHY SAP.

[0042] It should also be realized from the foregoing that the scheduling sublayer 24 implements a first-come, first-served scheduling discipline with respect to data requests at the point of contention. For example, referring to Figs. 8 and 9, VCs 1, 2, 3 and 4 submit SDU data requests in that order, so the scheduling sublayer 24 services the connections in that order. In the particular example shown, the SDU of VC 4 cannot be immediately transferred to the CS/SAR because no bandwidth is available due to the servicing of the other VCs and so the SDU of VC 4 is forwarded to the CS/SAR at time $t = 6$. Nevertheless, through the use of queues 34, 36 and 38, the scheduling sublayer 24 applies a round-robin service discipline with respect to the connections *per se*. For instance, referring to Figs. 8 and 9, suppose that at time $t = 0^+$, VC 1 presents a second SDU (i.e., two SDUs) to the scheduling sublayer 24 before VC 4 presents its SDU thereto. Due to step 42 of the preferred method, the scheduling sublayer 24 places the second SDU of VC 1 in VC₁ queue 38A. The SDU of VC 4 is placed in VC₄ queue 38D, but VC 4 is itself placed in the TX-pending queue 36. At time $t = 6$ when the first SDU of VC 1 has finished transmitting, step 64 of the preferred method notes that VC 1 has an outstanding data request and thus places VC 1 on the TX-pending queue 36, behind VC 4. Thus, at time $t = 6$ when bandwidth becomes available, the (only) SDU of VC 4 is submitted to the CS/SAR by step 86 ahead of the second SDU of VC 1, even though the latter physically arrived at the scheduling sublayer 24 prior to the former. In this manner, the preferred method provides a means for ensuring that no VC is starved.

[0043] In alternative embodiments, the scheduling sublayer could implement a strict first-come first-served order per SDU, e.g., by transferring each SDU to the CS/SAR in the order of arrival as follows: (i) immediately, provided that bandwidth is available, or in due course, when bandwidth becomes available to the VC of that SDU or (ii) immediately, if the VC of that SDU is currently transmitting an SDU, thereby essentially changing the queuing point for such VCs to the ATM Layer. Such an embodiment may, however, risk starving a connection.

[0044] The AAL according to preferred embodiment as shown in Fig. 5 places the scheduling sublayer 24 above the CS/SAR layers 12 and 13. It should be appreciated that in alternative embodiments the scheduling sublayer 24 may be placed between the SAR layer 12 and the ATM layer 11. In such an embodiment, the scheduling sublayer transfers groups of AAL SAR PDUs to the ATM layer 11, wherein each group of AAL SAR PDUs corresponds to an AAL SDU. This embodiment yields an AAL having the same functionality as that of the preferred embodiment because the payload of the group of AAL SAR PDUs is in effect the AAL SDU.

Service Classes Characterized by PCR and SCR

[0045] The embodiment described above operates well when the ATM virtual connections have traffic parameters such that they are always permitted to transmit at rates up to some peak cell rate. However, further optimization of the scheduling sublayer 24 can be obtained for the non-real-time variable bit rate (nrt-VBR) service category, where the rate at which a VC can transmit at a particular instant may vary between a specified sustainable cell rate and a peak cell rate.

[0046] ATM Forum document no. AF-TM-0056.000 and ITU specification no. I.371 describe the VBR traffic contract in terms of peak cell rate (PCR), sustainable cell rate (SCR) and maximum burst size (MBS). Compliance to the VBR traffic contract is defined in terms of the generic cell rate algorithm (GCRA), also commonly known as the "leaky bucket" algorithm, which is shown in Fig. 11. A VC is said to be compliant if the cell stream carried by that VC conforms to both GCRA(T_S, τ_S) and GCRA($T_p, 0$), where T_S is the cell period at the SCR,

$$(T_s = \frac{1}{SCR}).$$

5 T_p is the cell period at the PCR,

$$(T_p = \frac{1}{PCR}).$$

10

and τ_s is the burst tolerance, ($\tau_s = (MBS-1) * (T_s - T_p)$).

[0047] The essence of the non-real-time VBR traffic contract is that it explicitly allows a VC to transmit in a "bursty" fashion. That is, a nrt-VBR VC can transmit indefinitely at its SCR, and be in conformance with its traffic contract. Alternatively, the VC can cyclically transmit at rates above its SCR, up to its PCR, for a well-defined period, followed by another period during which it must transmit at rates below its SCR. This bursty transmission pattern is also in conformance with the nrt-VBR traffic contract.

[0048] Fig. 12 is a bandwidth occupancy chart illustrating a transmission pattern that arises under the scheduling sublayer 24 of the first preferred embodiment for two nrt-VBR VCs. Each of these VCs has a set of AAL SDUs to send, with traffic parameters specifying an SCR of 50% of the PHY bandwidth and a PCR of 100% of the PHY bandwidth. It will be seen that each VC is continually sending at its SCR, and there is no burstiness in transmission. This transmission pattern is in conformance with the example nrt-VBR traffic contract, and the average latencies are noted on the chart.

[0049] The second preferred embodiment improves upon the first preferred embodiment by modifying the scheduling sublayer 24 to schedule the transmission of an SDU such that the burstiness permitted by the nrt-VBR traffic contract is advantageously used to further reduce the average latency in the transmission of an SDU. This is done by scheduling the transmission of SDUs from various VCs such that the VCs may alternate between transmitting at rates approaching their PCRs and pausing their transmission, and by having the bursts in the transmission of each VC contain exactly one SDU. The transmission pattern under this approach is shown in Fig. 13 for the same example traffic contract described with reference to Fig. 12. It will be seen that while one VC has paused its transmission, the other VC is sending at its PCR. This transmission pattern is highly bursty, but also conforms to the example nrt-VBR traffic contract.

[0050] As in the first embodiment, the second embodiment employs ATM Layer equipment 11' which preferably has the following characteristics:

- Each VC whose cells are multiplexed by the ATM Layer may be assigned its own constant cell period. This cell period may be modified at any time by some other entity in the system, such as the scheduling sublayer 24, but it is not otherwise modified by the ATM layer.
- The ATM Layer performs cell multiplexing such that, for each opportunity it has to send a cell to the PHY, it may select a cell from any of the VCs for which it has candidate cells available for transmission, provided only that the interval between the current instant and the time at which the previous cell was transmitted for a particular VC is not less than the cell period specified for that VC.
- In the event that the cell period specified for a VC is modified by some entity such as the scheduling sublayer, the interval between the time at which the last cell was transmitted at the old cell period and the first cell is transmitted at the new cell period is not less than the lesser of the old and new cell periods.

45

[0051] Figs. 14A and 14B show the preferred method carried out by the scheduling sublayer 24 in accordance with the second preferred embodiment. Those steps which have been added or modified in relation to Figs. 6A and 6B of the first embodiment are shown in stippled outlines.

[0052] In the second preferred method, the scheduling sublayer 24 stores two parameters, t and X , for each VC. X corresponds to the "fill level" of the GCRA (T_s, τ_s) "leaky bucket", measured at time t . Initially, i.e., when a given VC is established, both X and t are zero. X and t are preferably recomputed by the scheduling sublayer 24 for each VC at two instants:

- 55 (i) at a step 47a (Fig. 14A), before the scheduling sublayer 24 passes (at step 52) a given SDU 22 to the CS/SAR 12, and
- (ii) at a step 61 (Fig. 14B), after the scheduling sublayer 24 receives (at event 60) an indication that the transmission of the given SDU 22 is complete.

[0053] At step 47a, before the scheduling sublayer 24 requests the SAR 12 to transmit a given SDU, X and t are calculated as follows:

$$X \leftarrow \max(X - t_{now} + t, 0) \quad (1)$$

$$t \leftarrow t_{now}$$

where t_{now} represents the time at which the calculation is performed. This calculation represents the "draining" of the leaky bucket in the period from t to t_{now} , during which there was no transmission of PDUs for the VC.

[0054] At step 61, after the scheduling sublayer 24 receives an indication that the transmission of the given SDU is complete, X and t are calculated as follows:

$$X \leftarrow \max(X - t_{now} + t + N * T_s, 0) \quad (2)$$

$$t \leftarrow t_{now}$$

where N is the number of PDUs (that is, the number of ATM cells) into which the given SDU was segmented by the SAR 12. This calculation represents the "filling" of the leaky bucket in the period from t to t_{now} , during which N PDUs were transmitted on the VC.

[0055] When the scheduling sublayer 24 has a candidate SDU which could begin transmission immediately, it checks whether the first cell of that SDU would be conformant with GCRA(T_s, τ_s) if the SDU were transmitted immediately, i.e., if $X - t_{now} + t \leq \tau_s$. Such a check occurs at step 47b (Fig. 14A), which corresponds to the condition where an AAL user sends a data request to the scheduling sublayer 24, PHY bandwidth is currently available, and queues 34 and 36 of pending SDUs for that VC are empty. This check also preferably occurs at step 81a (Fig. 14B) which corresponds to the situation where another SDU has completed transmission and the scheduling sublayer is scanning the TX-pending queue 36 of VCs with pending SDUs.

[0056] If the check, $X - t_{now} + t \leq \tau_s$, is satisfied, the scheduling sublayer 24 can pass the candidate SDU to the CS/SAR 12 for immediate transmission. If the check is not satisfied, the scheduling sublayer 24 cannot immediately begin transmitting the candidate SDU without violating GCRA(T_s, τ_s). Instead, in the case of a newly submitted SDU, at step 56' the SDU is placed at the end of the queue 38 of pending SDUs for that VC, and that VC is placed on the TX-pending queue 36 of VCs which have pending SDUs. In the case where the scheduling sublayer 24 is scanning the TX-pending queue 36 of VCs with pending SDUs, queue 36 is left untouched. In either case, at steps 47c (Fig. 14A) and 81b (Fig. 14B), the scheduling sublayer 24 starts a timer set to a period $X - \tau_s - t_{now} + t$. When this timer expires at event 79 (Fig. 14B), the scheduling sublayer will again scan the TX-pending queue 36 of VCs with pending data requests. At that time, the former candidate SDU should be able to be transmitted in conformance with GCRA(T_s, τ_s).

[0057] Before the scheduling sublayer passes an SDU to the CS/SAR for transmission (step 52), at step 48' the scheduling sublayer computes a cell period, T , (T being the inverse of the cell rate) which should be applied by the ATM Layer to the VC over the course of the transmission of the SDU.

[0058] When $N > 1$, T is calculated according to:

$$T \leftarrow \max\left(T_s + \frac{X - \tau_s}{N - 1}, T_p, T_L\right) \quad \text{if } N > 1 \quad (3a)$$

[0059] where T_L represents the period at which the remaining bandwidth of the PHY would be completely consumed by the transmission of this SDU. The first term of the above $\max()$ function represents the cell period which would cause the "bucket" of the VC to be completely filled at the end of the SDU's transmission. This implies that the SDU will be transmitted in as fast a burst as possible while still being in conformance with GCRA(T_s, τ_s). (Note that $X = \tau_s$ implies that the bucket is full when the SDU's transmission begins, in which case the first term of the $\max()$ function simplifies to T_s .) When the bucket is full, the cell period cannot be less than the sustainable cell period. The second term prevents the cell period from being less than the peak cell period. The third term prevents the VC from attempting to consume more PHY bandwidth than exists.

[0060] When $N = 1$, the first term in equation {3a} is not considered. Thus,

$$T \leftarrow \max(T_p, T_L) \quad (3b)$$

[0060] After an SDU's transmission is completed (event 60), the scheduling sublayer 24 attempts at step 70' to increase the cell rate, i.e., decrease the cell period, of any VCs with SDUs currently in transmission. To accommodate

5 this, the scheduling sublayer 24 stores parameters X_{est} , R_{est} and t_{est} for each VC. These represent the estimated bucket fill level, the estimated number of cells remaining to transmit in the SDU, and the time of estimation, respectively. They are initialized at step 52, when the scheduling sublayer 24 passes an SDU to the CS/SAR for transmission, according to:

$$X_{est} \leftarrow X, R_{est} \leftarrow N, t_{est} \leftarrow t \quad \{4\}$$

[0061] Thereafter, if the VC's cell period is adjusted while the SDU is in transmission, the new cell period is computed at step 70' according to:

$$X_{est} \leftarrow X_{est} + (t_{now} - t_{est}) \frac{T_s - T}{T} \quad \{5a\}$$

$$R_{est} \leftarrow R_{est} - \frac{t_{now} - t_{est}}{T} \quad \{5b\}$$

$$t_{est} \leftarrow t_{now} \quad \{5c\}$$

$$T \leftarrow \max\left(T_s + \frac{X_{est} - \tau_s}{R_{est} - 1}, T_p, T_L\right) \quad \text{if } R_{est} > 1 \quad \{5d\}$$

[0062] If $R_{est} \leq 1$, the cell period of the VC is not modified.

25 [0063] Again, the first term in the $\max()$ of equation {5d} results in a cell period T such that the bucket will be full when the SDU's transmission is complete.

[0064] The "est" subscript in the above expressions refer to the fact that X_{est} and R_{est} are estimates of the bucket fill level and number of cells remaining to be transmitted in the SDU, based on the elapsed time since the last estimate (or since the SDU's transmission was requested to the SAR) and the cell period; they are not based on direct indications 30 from the ATM layer or the SAR.

[0065] In the event that the ATM Layer can provide a direct indication of the number of cells remaining to be transmitted, then the scheduling sublayer need not track X_{est} separately from X . Alternatively, in the event that the ATM Layer does not provide direct indications at any time, including when the last cell of an AAL SDU has been transmitted, there is again no need to distinguish between X_{est} and X . X_{est} should be distinguished from X only if the calculations during 35 the transmission of an SDU are time based estimates and the calculations at the end of the transmission of an SDU are based on a direct indication from the ATM Layer. The advantage of this distinction is that any errors which may accumulate in the estimations are discarded when the direct indication is received at the ATM Layer.

[0066] The foregoing additions and modifications to the first preferred embodiment have the objective of determining the maximum possible cell rate which can be used for a particular VC at a particular instant in time while maintaining 40 conformance to the VBR traffic contract.

[0067] In addition to the foregoing additions and modifications, at steps 56' (Fig. 14A) and 66' (Fig. 14B) the scheduling sublayer 24 sorts the TX-pending queue 36 of VCs which have SDUs pending transmission according to increasing values of parameter t_E :

$$t_E \leftarrow X + t \quad \{6\}$$

[0068] t_E is simply the time at which the VC's bucket will be empty. This means that when one SDU's transmission has completed and another SDU's transmission may begin, the SDU selected to begin transmission will be the one on the VC whose bucket is emptiest. The VC with the emptiest bucket is that which can transmit at a cell rate closest to its 50 PCR, and is therefore the VC which can achieve the most bursty transmission.

[0069] Steps 56' and 66' attempt to maximize burstiness by taking advantage of the linearity of the GCRA. As shown in Fig. 15, there is a range of times when transmission of an SDU can begin which result in the same time for the completion of transmission of the last cell for that SDU, and hence the same latency therefor. Within that range, if the start of transmission of an AAL SDU is delayed such that the VC is not transmitting during that delay period, the VC's bucket 55 will be emptying. This allows the SDU to eventually be transmitted at a cell rate approaching the PCR, which recoups the delay time. Thus, the increase in burstiness has no impact on that latency of the SDU. However, in the idle time before that transmission of that SDU begins, it is possible that another SDU could be completely transmitted on another VC, thereby lowering the average SDU latency of the system.

[0070] An example is now presented of the advantages obtained by sorting the TX-pending queue 36 in steps 56' and 66' in order to maximize burstiness and hence reduce latency. Consider a system having three nrt-VBR VCs with each VC having the same traffic parameters: a PCR of 100% of the PHY bandwidth, an SCR of 50% of the PHY bandwidth, and an MBS of 4. This yields $T_P = 1$, $T_S = 2$ and $t_S = 3$ for each VC. (As before, a unit of time is deemed to be the time required to transmit one SAR PDU or cell at the PHY.)

[0071] Suppose at time 0, the scheduling sublayer 24 receives a data request on VC 1 for an SDU which will be segmented into 4 SAR PDUs, and immediately thereafter receives a data request on VC 2 for an SDU which will fit into a single SAR PDU. Just after time 4, the scheduling sublayer receives another data request on VC 1 for an SDU which will be segmented into 3 SAR PDUs. Also just after time 4, the scheduling sublayer receives a data request on VC 3 for an SDU which will be segmented into 3 SAR PDUs. Assume that all VCs are in their initial state, with empty buckets, at time 0.

[0072] Referring additionally to the bandwidth occupancy chart of Fig. 16, the scheduling sublayer proceeds as follows:

15 $t_{now}=0$: First, the scheduling sublayer verifies that VC 1 is GCRA (T_S, t_S), compliant, i.e., that $X \cdot t_{now} + t \leq t_S$ for VC 1. Since this condition is met, VC 1 can begin transmission of an AAL SDU immediately. The scheduling sublayer then computes $(X, t)=(0, 0)$ for VC 1. The period at which PDUs should be transmitted for this SDU is, from equation [3a], given by:

$$T = \max\left(T_S + \frac{X - t_S}{N - 1}, T_P, T_L\right) = \max\left(2 + \frac{0 - 3}{4 - 1}, 1, 1\right) = 1$$

20 In other words, the scheduling sublayer will tell the ATM Layer to set the cell rate of VC 1 to 100% of the PHY bandwidth, and will request the transmission of the SDU for VC 1 to the SAR.

25 Also, since all the PHY bandwidth has been allocated at this time, the SDU requested for VC 2 is added to the end of the VC₂ queue 38b of pending SDUs for VC 2, and VC 2 is placed on the TX-pending queue 36 for VCs having pending data requests, with $t_E=0$.

30 $t_{now}=4$: The transmission of the SDU for VC 1 is completed. The parameters $(X, t) = (0 - 4 + 0 + 2 * 4, 4) = (4, 4)$ are stored against VC 1. The transmission latency for this SDU is 4. There are no other VCs with SDUs in transmission. The TX-pending queue 36 of VCs having pending data requests is scanned, finding VC 2. For VC 2, the condition $X \cdot t_{now} + t \leq t_S$ is met, so transmission may begin. The parameters $(X, t)=(0, 4)$ are computed for VC 2. Since $N=1$ for the SDU requested on VC 2, the transmission of the SDU occurs at the peak cell period of 1. The scheduling sublayer 24 thus signals the ATM Layer 11 to set the cell rate of VC 2 to 100% of the PHY bandwidth, and will request the transmission of the SDU for VC 2 to the SAR 12.

35 $t_{now}=4^+$: Next, the scheduling sublayer receives the data request for the second SDU on VC 1. Since there is no unused PHY bandwidth remaining, the SDU requested for VC 1 is added to the end of VC₁ queue 38a of pending SDUs for VC 1, and VC 1 is placed on the TX-pending queue 36 for VCs with pending SDU data requests, with $t_E=X+t=8$.

40 Finally, the scheduling sublayer receives the data request for the SDU of VC 3. Since there is no unused PHY bandwidth remaining, the SDU for VC 3 is added to the end of the VC₃ queue 38c of pending SDUs for VC 3, and VC 3 is placed on the TX-pending queue 36 of VCs with pending SDUs. Since the t_E value for VC 3 is 0, which is lower than the t_E value for VC 1 ($t_E = 8$ for VC 1), VC 3 is inserted before VC 1 on the TX-pending queue 36 for VCs having SDU data requests pending, in accordance with step 66'.

45 $t_{now}=5$: The transmission of the SDU for VC 2 is completed. The parameters $(X, t) = (0 - 5 + 4 + 2 * 1, 5) = (1, 5)$ are stored against VC 2. The transmission latency for this SDU is 2. There are no other VCs with SDUs in transmission. The TX-pending queue 36 of VCs having pending SDUs is scanned, finding VC 3 first. For VC 3, the condition $X \cdot t_{now} + t \leq t_S$ is met, so transmission may begin. The parameters $(X, t)=(0, 5)$ are computed for VC 3. The period at which PDUs should be transmitted for the SDU of VC 3 is given by:

$$T = \max\left(T_S + \frac{X - t_S}{N - 1}, T_P, T_L\right) = \max\left(2 + \frac{0 - 3}{3 - 1}, 1, 1\right) = 1$$

EP 0 952 753 A2

The scheduling sublayer thus signals the ATM layer to set the cell rate of VC 3 to 100% of the PHY bandwidth, and will request the transmission of the SDU for VC 3 to the SAR.

- 5 $t_{now}=8$: The transmission of the SDU for VC 3 is completed. The parameters $(X, t) = (0 - 8 + 5 + 2 * 3, 8) = (3, 8)$ are stored against VC 3. The transmission latency for this SDU is 4.
 There are no other VCs with SDUs in transmission. The TX-pending queue 36 of VCs having pending SDU data requests is scanned, finding VC 1. For VC 1, the condition $X-t_{now}+t \leq \tau_s$ is met, so transmission may begin. The parameters $(X, t) = (4 - 8 + 4, 8) = (0, 8)$ are computed for VC 1. The period at which PDUs should be transmitted for the SDU of VC 3 is given by:

10

$$T = \max\left(T_s + \frac{X - \tau_s}{N - 1}, T_p, T_L\right) = \max\left(2 + \frac{0 - 3}{3 - 1}, 1, 1\right) = 1$$

15 The scheduling sublayer thus signals the ATM Layer to set the cell rate of VC 1 to 100% of the PHY bandwidth, and will request the transmission of the second SDU for VC 1 to the SAR.

- 20 $t_{now}=11$: The transmission of the second SDU for VC 1 is completed. The parameters $(X, t) = (0 - 11 + 8 + 2 * 3, 11) = (3, 11)$ are stored against VC 3. The transmission latency for this SDU is 7.
 The average transmission latency per unit of SDU data is then $(4 * 4 + 2 * 1 + 4 * 3 + 7 * 3) / (4 + 1 + 3 + 3) \approx 4.6$.

25 [0073] Referring now to the bandwidth occupancy chart of Fig. 17, a contrasting example is now presented assuming the sorting steps 56' and 66' of the second preferred embodiment are not applied so that the TX-pending queue 36 of VCs with SDUs pending is maintained in simple first-in-first-out (FIFO) order as in the first preferred embodiment. In this case, after $t_{now}=4$, the order of the TX-pending queue 36 of VCs having pending SDUs would have VC 1 before VC 3. Accordingly, the following modifications would occur:

- 30 $t_{now}=5$: The transmission of the SDU for VC 2 is completed. The parameters $(X, t) = (0 - 5 + 4 + 2 * 1, 5) = (1, 5)$ are stored against VC 2. The transmission latency for this SDU is 2.
 There are no other VCs with SDUs in transmission. The TX-pending queue 36 of VCs with pending SDUs is scanned, finding VC 1 first. For VC 1, the condition $X-t_{now}+t \leq \tau_s$ is met, so transmission may begin. The parameters $(X, t) = (4 - 5 + 4, 5) = (3, 5)$ are computed for VC 1. The period at which PDUs should be transmitted for the second SDU of VC 1 is given by:

$$T = \max\left(T_s + \frac{X - \tau_s}{N - 1}, T_p, T_L\right) = \max\left(2 + \frac{3 - 3}{3 - 1}, 1, 1\right) = 2$$

- 40 The scheduling sublayer thus instructs the ATM Layer to set the cell rate of VC 1 to 50% of the PHY bandwidth, and will request the transmission of the second SDU for VC 1 to the SAR.
 Fifty per cent (50%) of the PHY bandwidth remains unused, so the scanning of the TX-pending queue 36 of VCs with SDUs pending continues. VC 3 is found. For VC 3, the condition $X-t_{now}+t \leq \tau_s$ is met, so transmission may begin. The parameters $(X, t) = (0, 5)$ are computed for VC 3. The period at which PDUs should be transmitted for VC 3's SDU is given by:

45

$$T = \max\left(T_s + \frac{X - \tau_s}{N - 1}, T_p, T_L\right) = \max\left(2 + \frac{0 - 3}{4 - 1}, 1, 2\right) = 2$$

- 50 This time, the limitation of the bandwidth allocated to VC 3 is not imposed by GCRA conformance, but by the limited amount of PHY bandwidth remaining. The scheduling sublayer thus instructs the ATM Layer to set the cell rate of VC 3 to 50% of the PHY bandwidth, and will request the transmission of the SDU for VC 3 to the SAR.

- 55 $t_{now}=11$: The transmission of the second SDU for VC 1 and the SDU for VC 3 are both completed. The parameters $(X, t) = (3 - 11 + 5 + 2 * 3, 11) = (3, 11)$ are stored against VC 1; the parameters

EP 0 952 753 A2

$(X, t) = (0 - 11 + 5 + 2 * 3, 11) = (0, 11)$ are stored against VC 3. The transmission latency for both SDUs is 7.

The average transmission latency per unit of SDU data is then $(4 * 4 + 2 * 1 + 7 * 3 + 7 * 3) / (4 + 1 + 3 + 3) \approx 5.5$ - worse than is achieved when the queue of VCs with SDUs pending is sorted by increasing t_E .

[0074] The discussion now turns to mathematically demonstrating that the second preferred embodiment produces a cell stream conformant to the nrt-VBR traffic contract. The analysis refers to the algorithms and discussion presented in ITU specification I.371 and especially in ATM Forum document AF-TM-0056.000, Appendix C, and uses a modified version of the notation in the "continuous-state leaky bucket algorithm" presented as an equivalent of the GCRA in AF-TM-0056.000 and I.371.

[0075] Let t_K be the time at which the K^{th} conforming cell was transmitted on an ATM virtual connection, let X'_K be the trial value of the "leaky bucket counter" before the K^{th} conforming cell was transmitted, and let X_K be the value of the "leaky bucket counter" of the GCRA(T_S, τ_S) after the K^{th} conforming cell was transmitted. If a conforming cell is transmitted at t_K , then the relationship between X'_K and X_K is given by:

$$X_K = X'_K + T_S \quad \{7\}$$

[0076] Suppose that at time t'_{K+1} , the SAR requests of the ATM layer that N ATM SDUs, that is, cells, be transmitted. The minimum values of t_{K+1} and t_{K+N} — that is, the earliest time at which the first and last of the N cells can be transmitted — can be calculated such that compliance to the traffic contract is guaranteed. Further, it can be shown that if $N > 1$, the N cells can be transmitted at some rate with constant period T. The linearity of the GCRA guarantees that if the first and last of N cells are conforming, and all cells between the first and last are equally spaced, then all cells between the first and last will also be conforming.

[0077] Note that transmitting cells at a fixed rate with period T is but one of many possible transmission scenarios that would be compliant to a given traffic contract.

[0078] First, the minimum value of t_{K+1} is calculated:

$$t_{K+1} \geq \max(t'_{K+1}, t_K + T_P, t_K + X_K - \tau_S) \quad \{8\}$$

[0079] The second term of the max() function in {8} assures compliance with GCRA($T_P, 0$); the last term assures compliance with GCRA(T_S, τ_S). The first term is pro forma, and indicates that the ATM layer cannot send the first cell before the ATM layer user has requested it.

[0080] Before the first cell is transmitted at t_{K+1} , the value of X'_{K+1} is given by:

$$X'_{K+1} = \max(X_K - t_{K+1} + t_K, 0) \quad \{9\}$$

The first term of the max() function in {9} corresponds to the case where there has not been sufficient time between t_K and t_{K+1} to "empty the bucket"; the second (zero) term of the max() function indicates that there has been sufficient time between t_K and t_{K+1} to "empty the bucket".

[0081] Next, if $N > 1$, assume there will be $N - 1$ iterations through the GCRA after t_{K+1} up to and including t_{K+N} , where each iteration is separated by equal time T. After these $N - 1$ iterations, to conform to GCRA(T_S, τ_S), T must satisfy:

$$(N-1)(T_S - T) + X'_{K+1} \leq \tau_S \quad \{10\}$$

[0082] To conform to GCRA($T_P, 0$), T must satisfy $T = T_P$. Therefore, combining these constraints on T, obtain for the minimum value of T:

$$T \geq \max\left(T_S + \frac{X'_{K+1} - \tau_S}{N-1}, T_P\right) \quad \{11\}$$

[0083] Similarly, after $N - 1$ iterations through the GCRA, calculate X_{K+N} , according to:

$$X_{K+N} = NT_S - (N-1)T + X'_{K+1} \quad \{12\}$$

[0084] Note that {12} applies for all N, including $N = 1$. For $N = 1$, however, the need to calculate T is obviated by the multiplication of T by $N - 1$.

[0085] When this algorithm is first run at time t'_1 for a particular VC, the values t_0 and X_0 are initialized to

$$t_0 = t'$$

5 and $X_0 = 0$.

[0086] It can be shown that relaxing the condition that the N cells be transmitted with equal period T does not change the minimum time t_{K+N} , but the loss of constant T does cause the loss of the guarantee that cells between the first and Nth are conformant. Therefore, there is no benefit but there is some cost in spacing the N cells equally.

10 [0087] Referring back to the operation of the second preferred embodiment, the scheduling sublayer's per-VC parameter X immediately before the scheduling sublayer requests transmission of an AAL SDU of the SAR, corresponds to X'_{K+1} in this analysis; the parameter t corresponds to t_{K+1} in the analysis. Once the ATM layer has completed transmission of all SAR PDUs and has indicated this back to the scheduling sublayer, the scheduling sublayer's per-VC parameter X corresponds to X_{K+N} in this analysis; the parameter t corresponds to t_{K+N} in the analysis. Thus, {9} gives {1}, {12} gives {2}, {11} gives {3a}, and {8} gives the conditions stated in the second desired characteristic of the ATM layer equipment.

15 [0088] The minimum time at which the last conformant SAR PDU of an AAL SDU can be transmitted by the ATM layer is then t_{K+N} , where:

$$\begin{aligned} 20 \quad t_{K+N} &= t_{K+1} + (N - 1)T \\ &= t_{K+1} + (N - 1)\max\left(T_S + \frac{X'_{K+1} - \tau_S}{N - 1}, T_P\right) \\ &= t_{K+1} + (N - 1)\max\left(T_S + \frac{\max(X_K - t_{K+1} + t_K, 0) - \tau_S}{N - 1}, T_P\right) \end{aligned} \quad \{13\}$$

25

[0089] The expansions in {13} follow from {11} and {9}. Now, suppose the following inequalities are satisfied:

$$30 \quad X_K - t_{K+1} + t_K \geq 0 \quad \{14a\}$$

$$T_S + \frac{X'_{K+1} - \tau_S}{N - 1} \geq T_P \quad \{14b\}$$

35 [0090] Then, {13} can be simplified to:

$$\begin{aligned} 40 \quad t_{K+N} &= t_{K+1} + (N - 1)\left(T_S + \frac{S_K - t_{K+1} + t_K - \tau_S}{N - 1}\right) \\ &= (N - 1)T_S + X_K + t_K - \tau_S \end{aligned} \quad \{15\}$$

[0091] The result of {15} means that t_{K+N} is independent of t_{K+1} provided conditions {14a} and {14b} are met. The inequalities {8}, {14a} and {14b} can be combined with the definition of T_S and rearranged as:

$$45 \quad X_K + t_K - \tau_S \leq t_{K+1} \leq X_K + t_K + \min((N - MBS)(T_S - T_P), 0) \quad \{16\}$$

[0092] This gives the range over which the time of transmission of the first SAR PDU for an AAL SDU may vary, without affecting the time of transmission of the last SAR PDU for the AAL SDU and hence the latency of the AAL SDU. In 50 other words, this is the time range illustrated in Fig. 15.

[0093] Those skilled in the art will appreciate that numerous modifications and variations may be made to the preferred embodiments without departing from the spirit and scope of the invention.

Claims

55

1. Apparatus for transmitting messages associated with a plurality of variable bit rate connections, each of which is associated with a traffic contract which defines compliance thereto as conformance to a leaky bucket algorithm, said apparatus comprising:

transmission equipment for receiving multiple messages, segmenting each received message into one or more data units, and multiplexing such data units from various connections into a single output stream for transport over a physical interface at an output transmission rate;

bandwidth provisioning means associated with the transmission equipment for dynamically providing a portion of the output transmission rate to any connection; and

a scheduler for scheduling the transfer of messages to the single stream of the transmission equipment and for directing the bandwidth provisioning means to allocate a portion of the output transmission rate to each connection at the time its message is transferred to the transmission equipment, said portion of the output transmission rate for a given connection being substantially equal to $1/T$, T being computed as

$$T \leftarrow \max\left(T_S + \frac{X - \tau_S}{N - 1}, T_p, T_L\right), \quad \text{if } N > 1$$

and

$$T \leftarrow \max(T_p, T_L), \quad \text{if } N = 1$$

where T_S is a period corresponding to a constant sustained transmission rate, T_p is a period corresponding to a peak transmission rate, τ_S is a burst tolerance, N is the number of data units in the message, X is a fill level of the leaky bucket associated with the given connection, and T_L corresponds to all unused bandwidth allocated to a service class to which the given connection belongs.

2. Apparatus for transmitting messages associated with a plurality of variable bit rate connections, each of which is associated with a traffic contract which defines compliance thereto as conformance to a leaky bucket algorithm, said apparatus comprising:

transmission equipment for receiving multiple messages, segmenting each received message into one or more data units, and multiplexing such data units from various connections into a single output stream for transport over a physical interface at an output transmission rate;

bandwidth provisioning means associated with the transmission equipment for dynamically providing a portion of the output transmission rate to any connection; and

a scheduler, for scheduling the transfer of messages to the single stream of the transmission equipment, the scheduler being enabled to transfer the messages according to an order corresponding to the level of a leaky bucket associated with each corresponding connection in order to reduce the interleaving of data units from various connections over said single stream.

3. The apparatus according to claim 1 or claim 2, wherein:

the transmission equipment comprises a segmentation means and a multiplexing means for respectively segmenting each received message into said data units and multiplexing said data units into said single stream; and

the scheduler transfers messages to said single stream by one of

(a) receiving the messages directly from the connections and transferring the messages to the segmentation means which substantially immediately transfers the messages to the multiplexing means, and

(b) receiving the data units from the segmentation means and transferring groups of data units to the multiplexing means, wherein each such group corresponds to a message.

4. The apparatus according to claim 3, wherein the scheduler is enabled to set a new transmission rate for a given connection during the time a message associated therewith is in the process of being transmitted by the transmission equipment, said new transmission rate being substantially equal to $1/T$, T being computed as

$$T \leftarrow \max\left(T_S + \frac{X_{est} - \tau_S}{R_{est} - 1}, T_p, T_L\right), \quad \text{if } R_{est} > 1,$$

where X_{est} and R_{est} respectively represent an expected bucket fill level and an expected number of data units remaining to be transmitted in the message for the given connection.

5. The apparatus according to claim 3, wherein the scheduler maintains a list of connections having pending data requests and buffers non-immediately serviceable messages in connection buffers.
6. The apparatus according to claim 5, wherein the scheduler sorts the list of connections having pending data requests in an order corresponding to respective instants of time when the leaky buckets associated with the listed connections will substantially become empty.
- 10 7. The apparatus according to claim 5, wherein the scheduler, in processing a message, enqueues the message in the appropriate connection buffer in the event the associated connection currently has a message being transmitted by the transmission equipment or the associated connection is in the list of connections having pending data requests.
- 15 8. The apparatus according to claim 5, wherein the scheduler is enabled to process a message by placing its associated connection in the list of connections having pending data requests and enqueueing the message in the appropriate connection buffer in the event no bandwidth is available to the associated connection.
- 20 9. The apparatus according to claim 5, wherein the scheduler is enabled to process a message by placing its associated connection in the list of connections having pending data requests and enqueueing the message in the appropriate connection buffer in the event the associated connection is not GCRA(T_S, τ_S) conformant.
- 25 10. The apparatus according to claim 9, wherein, in event a connection is not GCRA(T_S, τ_S) compliant such that the associated message is enqueued in the corresponding connection buffer, the scheduler schedules the transfer of the associated message to said single stream for a time when the connection will be GCRA(T_S, τ_S) conformant.
- 30 11. The apparatus according to claim 10, further including indication means for informing the scheduler when a message has been completely transmitted by the transmission equipment, and wherein upon the actuation of the indication means in respect of a given connection, the scheduler is enabled to check the corresponding connection buffer connection for the existence of pending messages and in such event to place the given connection on the list of connections having pending data requests.
- 35 12. The apparatus according to claim 11, wherein the scheduler sorts the list of connections having pending data requests in an order corresponding to respective instants of time when the leaky buckets associated with the listed connections will substantially become empty.
- 40 13. A method for transmitting multiple messages associated with multiple variable bit rate connections over a network, each such connection being associated with a traffic contract which defines compliance thereto as conformance with a leaky bucket algorithm, said method comprising:

segmenting each message of each connection into one or more transportable data units;
multiplexing the data units from various connections into a single stream for transport over a transmission media to a destination;
scheduling the transfer of messages from the multiple connections to the multiplexing step and;
allocating a portion of the transmission media bandwidth to each connection substantially when its message is transferred to the multiplexing step, said portion of the transmission media bandwidth for a given connection being substantially equal to $1/T$, T being computed as

$$50 \quad T \leftarrow \max\left(T_S + \frac{X - \tau_S}{N - 1}, T_p, T_L\right), \quad \text{if } N > 1$$

and

$$55 \quad T \leftarrow \max(T_p, T_L), \quad \text{if } N = 1$$

where T_S is a period corresponding to a constant sustained transmission rate, T_p is a period corresponding to

a peak transmission rate, τ_S is a burst tolerance, N is the number of data units in the message, X is a fill level of the leaky bucket associated with the given connection, and T_L corresponds to all unused bandwidth allocated to a service class to which the given connection belongs.

- 5 14. A method for transmitting multiple messages associated with multiple variable bit rate connections over a network, each such connection being associated with a traffic contract which defines compliance thereto as conformance with a leaky bucket algorithm, said method comprising:

10 segmenting each message of each connection into transportable data units; multiplexing the data units from various connections into a single stream for transport over a transmission media to a destination; and scheduling the transfer of messages from the multiple connections to the multiplexing step according to an order corresponding to the level of a leaky bucket associated with each corresponding connection in order to reduce the interleaving of data units from various connections over said single stream.

- 15 15. The method according to claim 13 or claim 14, wherein the scheduling step comprises one of:

20 (i) scheduling the transfer of groups of data units from the segmentation step to the multiplexing step, wherein each such group of data units corresponds to a message, and

25 (ii) scheduling the transfer of messages from the multiple connections to the segmentation step, wherein the segmentation step substantially immediately transfers the data units to the multiplexing step.

- 25 16. The method according to claim 15, wherein the scheduling step is enabled to schedule the transfer of messages to said single stream according to an ordering of connections having outstanding data requests, said ordering corresponding to respective instants of time when the leaky buckets associated with such connections will substantially become empty.

- 30 17. The method according to claim 16, wherein the scheduling step is enabled to set a new portion of the transmission media bandwidth for a given connection during the time a message associated therewith is in the process of being transmitted by the multiplexing step, said new bandwidth portion being substantially equal to $1/T$, T being computed as

$$35 T \leftarrow \max\left(T_S + \frac{X_{est} - \tau_S}{R_{est} - 1}, T_p, T_L\right), \quad \text{if } R_{est} > 1,$$

40 where X_{est} and R_{est} respectively represent an expected bucket fill level and an expected number of data units remaining to be transmitted in the message.

- 45 18. The method according to claim 17, including the step of maintaining a list of connections having pending data requests and buffering non-immediately serviceable messages in connection buffers, wherein the list of connections having pending data requests is sorted in an order corresponding to respective points in time when the leaky buckets associated with the listed connections will substantially become empty.

- 45 19. The method according to claim 18, wherein the scheduling step is enabled to process a message by placing its associated connection in the list of connections having pending data requests and enqueueing the message in the appropriate connection buffer in the event the associated connection is not $GCRA(T_S, \tau_S)$ conformant.

- 50 20. Apparatus for transmitting ATM adaption layer (AAL) service data units (SDUs) associated with a plurality of variable bit rate connections, each of which is associated with a traffic contract which defines compliance thereto as conformance to $GCRA(T_S, \tau_S)$ and $GCRA(T_p, 0)$, said apparatus comprising:

55 segmentation means for segmenting each AAL SDU into at least one AAL segmentation and reassembly protocol data units (SAR PDUs);

ATM layer means for multiplexing the SAR PDUs of various connections into a single stream having an output transmission rate, said ATM layer means including bandwidth provisioning means for dynamically providing a portion of the output transmission rate to any connection; and

physical transmission means, connected to the ATM layer means, for transporting the single stream of SAR PDUs over a physical interface; and

a scheduler for scheduling the transfer of SDUs to the ATM layer means and for directing the bandwidth provisioning means to allocate a portion of the output transmission rate to each connection at the time its SDU is transferred to the ATM Layer means, said portion of the output transmission rate for a given connection being substantially equal to $1/T$, T being computed as

$$T \leftarrow \max\left(T_S + \frac{X - \tau_S}{N - 1}, T_p, T_L\right), \quad \text{if } N > 1$$

10

and

$$T \leftarrow \max(T_p, T_L), \quad \text{if } N = 1$$

15

where T_S is a period corresponding to a constant sustained transmission rate, T_p is a period corresponding to a peak transmission rate, τ_S is a burst tolerance, N is the number of data units in the message, X is a fill level of a leaky bucket associated with the given connection, and T_L corresponds to all unused bandwidth allocated to a service class to which the given connection belongs.

20

21. Apparatus for transmitting ATM adaption layer (AAL) service data units (SDUs) associated with a plurality of variable bit rate connections, each of which is associated with a traffic contract which defines compliance thereto as conformance to a leaky bucket algorithm, said apparatus comprising:

25

segmentation means for segmenting each AAL SDU into at least one AAL segmentation and reassembly protocol data units (SAR PDUs);

ATM layer means for multiplexing the SAR PDUs of various connections into a single stream having an output transmission rate, said ATM layer means including bandwidth provisioning means for dynamically providing a portion of the output transmission rate to any connection;

30

physical transmission means, connected to the ATM layer means, for transporting the single stream of SAR PDUs over a physical interface;

a scheduler for scheduling the transfer of SDUs to the ATM layer means, the scheduler being enabled to transfer the SDUs according to an order corresponding to the level of a leaky bucket associated with each corresponding connection in order to reduce the interleaving of data units from various connections over said single stream.

35

22. The apparatus according to claim 20 or claim 21, wherein the scheduler transfers SDUs to the ATM layer means by one of:

40

(a) receiving the SDUs from AAL users and transferring the SDUs to the segmentation means, wherein the segmentation means substantially immediately transfers the SAR PDUs to the ATM layer means; and

(b) transferring groups of SAR PDUs to the ATM layer means, wherein each such group corresponds to an SDU.

45

23. The apparatus according to claim 22, wherein the scheduler is enabled to schedule the transfer of SDUs to the ATM layer means according to an ordering of connections having outstanding data requests, said ordering corresponding to respective instants of time when the leaky buckets associated with such connections will substantially become empty.

50

24. The apparatus according to claim 23, wherein the scheduler is enabled to set a new transmission rate for a given connection during the time an SDU associated therewith is in the process of being transmitted by the ATM layer means, said new transmission rate being substantially equal to $1/T$, T being computed as

55

$$T \leftarrow \max\left(T_S + \frac{X_{est} - \tau_S}{R_{est} - 1}, T_p, T_L\right), \quad \text{if } R_{est} > 1.$$

EP 0 952 753 A2

where X_{est} and R_{est} respectively represent an expected bucket fill level and an expected number of SAR PDUs remaining to be transmitted in the SDU for the given connection.

5

10

15

20

25

30

35

40

45

50

55




Fig. 1 -





Fig. 2



PRIOR ART

Fig. 3




Fig. 4A




Fig. 4B




Fig. 5




Fig. 6A








Fig. 6B

**Fig. 7**

VC	SDU request time	Number of PDUs in SDU	PCR	Round-robin latency	Scheduling sublayer latency
1	0	3	50%	10.5	6
2	0	2	33%	8	6
3	0	4	25%	16	18
4	0	3	75%	10.5	10
Average latency				11.9	11

Fig. 8**Fig. 9****Fig. 10**



VIRTUAL SCHEDULING ALGORITHM

TAT Theoretical Arrival Time

$t_{a(k)}$ Time of arrival of a cell

I Increment

L Limit

At the time of arrival t_a of the first cell of the connection, $TAT = t_a(I)$

CONTINUOUS-STATE LEAKY BUCKET ALGORITHM

X Value of the Leaky Bucket counter

X' auxiliary variable

LCT Last Compliance Time

At the time of arrival t_a of the first cell of the connection, $X = 0$ and $LCT = t_a(k)$

Fig. 11




Fig. 12




Fig. 13




Fig. 14A




Fig. 14B




Fig. 15




Fig. 16




Fig. 17



(12)

EUROPEAN PATENT APPLICATION

(88) Date of publication A3:
 27.12.2000 Bulletin 2000/52

(51) Int. Cl.⁷: H04L 12/56, H04Q 11/04

(43) Date of publication A2:
 27.10.1999 Bulletin 1999/43

(21) Application number: 99106051.8

(22) Date of filing: 25.03.1999

(84) Designated Contracting States:
 AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
 MC NL PT SE

Designated Extension States:
 AL LT LV MK RO SI

(30) Priority: 30.03.1998 CA 2233491
 11.12.1998 US 208984

(72) Inventor: Dubuc, Ken
 Ottawa, Ontario K2B 8K3 (CA)

(74) Representative:
 Prato, Roberto et al
 STUDIO TORTA S.r.l.,
 Via Vlotti 9
 10121 Torino (IT)

(71) Applicant:
 Newbridge Networks Corporation
 Kanata, Ontario K2K 2E6 (CA)

(54) ATM adaption layer traffic scheduling for variable bit rate connections

(57) The ATM apparatus for transmitting messages associated with a plurality of variable bit rate connections comprises ATM layer transmission equipment for receiving multiple messages, segmenting each received message into a plurality of smaller data units, such AAL SAR PDUs, and multiplexing such data units into a single stream for transport over a physical interface. A scheduler receives messages from each of the variable bit rate connections and transfers the messages to the transmission equipment in an order corresponding to the level of a leaky bucket associated with each connection. The scheduler also dynamically sets a transmission rate, 1/T, for each connection at the time its message is transferred to the transmission equipment, where

$$T \leftarrow \max\left(T_S + \frac{X - \tau_S}{N - 1}, T_p, T_L\right), \quad \text{if } N > 1$$

and

$$T \leftarrow \max(T_p, T_L), \quad \text{if } N = 1$$

where T_S is a period corresponding to a constant sustained transmission rate, T_p is a period corresponding to a peak transmission rate, τ_S is a burst tolerance, N is the number of data units in the message, X is a fill level of the leaky bucket associated with the given connection, and T_L corresponds to all unused bandwidth allocated to a service class to which the given connection

belongs. These actions increase the burstiness of the connections and minimize interleaving of data units from the various connections in order to reduce the average latency (per unit of message) in transmitting the messages across a network.




Fig. 1




Fig. 5



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number

EP 99 10 6051

DOCUMENTS CONSIDERED TO BE RELEVANT					
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.)		
A	<p>TRANCHIER D P ET AL: "FAST BANDWIDTH ALLOCATION IN ATM NETWORKS" PROCEEDINGS OF THE INTERNATIONAL SWITCHING SYMPOSIUM, JP, TOKYO, IEICE, vol. SYMP. 14, 25 October 1992 (1992-10-25), pages 7-11, XP000337692 * page 7, left-hand column, line 35 - line 44 *</p> <p>XIE G G ET AL: "REAL-TIME BLOCK TRANSFER UNDER A LINK-SHARING HIERARCHY" IEEE / ACM TRANSACTIONS ON NETWORKING, US, IEEE INC. NEW YORK, vol. 6, no. 1, 1 February 1998 (1998-02-01), pages 30-41, XP000733549 ISSN: 1063-6692 * page 31, left-hand column, line 39 - line 46 *</p>	1-24	HO4L12/56 HO4Q11/04		
			TECHNICAL FIELDS SEARCHED (Int.Cl.)		
			HO4L HO4Q		
The present search report has been drawn up for all claims					
Place of search	Date of completion of the search	Examiner			
THE HAGUE	23 October 2000	Veen, G			
CATEGORY OF CITED DOCUMENTS					
X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background C : non-written disclosure P : intermediate document					
T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document					